

19

Intelligence Revisited

"I think, therefore I am."

Rene Descartes

Much of the motivation for this chapter came from a series of e-mail discussions I had with Julie Marble, a cognitive psychologist involved in robotics at the Idaho National Laboratory in Idaho Falls. Our respective organizations had been working for decades (independently at first, then cooperatively as part of a strategic alliance) towards enabling increasingly sophisticated robotic behaviors, with progressively more focus on human-robot teaming (Pacis, et al., 2004). By 2004 it seemed we had achieved enough truly useful autonomous functionality to where a major reassessment of roles and interactions was clearly in order: the robot was at an evolutionary turning point where it could soon become as big a part of the team as the human, instead of just an assistive subordinate tool to be used in elective fashion. This realization on our part put things in a whole new perspective.

The previously clear-cut distinctions between human and robot were beginning to blur, just as the boundary between *reality* and *virtuality* had already blurred (Everett, et al., 2004). The hardcore futurists were serving up unsettling (and in some cases even shocking) visions of where all this would ultimately lead, some painting scenarios where we humans would meld with machines, abandoning our frail bodies and transferring our brains into more capable computers (Moravec, 1988, 1999). Others took the view that computers, and hence robots – once they became intellectually superior to humans – would have no further use for us on what would then be their planet (Kurzweil, 1999). An even more apocalyptic speculation has the machine life-form coldly farming selected remnants of the human race as a ready source of slave labor (Warwick, 1997). All this wild and thought-provoking stuff was delivered not by over-hyped science-fiction writers, but coming instead from some of the most intelligent, respected, and experienced experts in the field itself.

In this chapter we begin to explore the theoretical upper limit of robotic intelligence, in terms of a machine's ability to actually think like a human being. Interestingly

The Evolution of ROBART

enough, this issue is basically a natural extension of the very question I posed in the opening pages of my thesis, almost 25 years ago (Everett, 1982):

“There exists at the one end simply a pre-programmed dedicated controller able to repeatedly execute the most complex of instructions and effect the motion of actuators, valve positions, motor speeds, etc. On the other end of the spectrum, however, there are evolving machines that can function on their own, evaluating their changing environment, and reacting as needed to carry out their intended tasks with no human intervention in such a way that they truly appear ‘intelligent.’ The obvious question arising is ‘At what point do these machines become robots?’”

In 1982 I was addressing the fuzzy distinction between automated machinery and intelligent robots, whereas in the following pages we contemplate the possibility of *artificial conscious awareness* on the part of “spiritual” machines (Figure 19-1). In so doing, I’m only lightly touching on a very complex and controversial topic that in its own right should be the subject of an entire book, which I’ll be the first to admit I’m not best qualified to write. So take it in that context, and I will attempt to make amends for my presumptuous intrusion into such territory by providing ample references to those far more experienced.

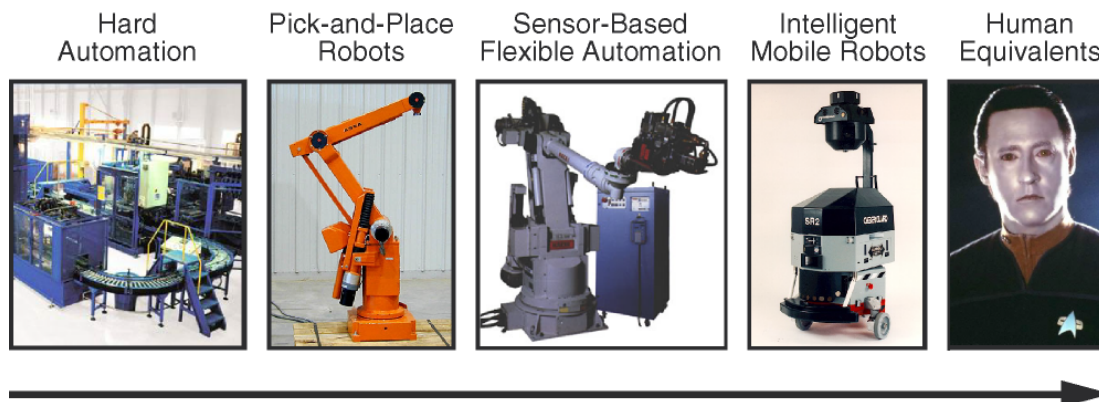


Figure 19-1. Generalized timeline of robotic evolution, from basic factory automation to the possibility of human-equivalent machines.

19.1 The Meaning of Intelligence

To better understand the nuances of the terminology *artificial intelligence*, it seems only prudent to start with precisely that which we seek to artificially replicate, namely *natural* (i.e., biological, primarily animal and human) *intelligence*. Unfortunately, however, the subject of *natural intelligence* is not all that well understood either.

19.1.1 Natural Intelligence

Researchers are clearly divided on the underlying issues that determine one’s general level of intellectual ability. One camp believes that intelligence is biologically constrained, and thus our ultimate mental potential is principally dictated by the genes we inherit from our parents, for better or worse. Others argue that social factors are the more dominant influence; our surrounding environment (i.e., culture, class, family, educational

Chapter 23 Intelligence Revisited

opportunities, gender) shapes our intellect, and there are no fundamental differences inherited at birth to help or hinder our progress. Meanwhile a third option, to which I myself subscribe, is that some combination of both these theories (i.e., of nature versus nurture) is more likely the case: we are born with some general intellectual capacity that is a function of our genes, and where we actually wind up across the full spectrum of human possibilities is determined by how these genes interact with our environment.

When it comes to quantifying intelligence, the general population tends to think in terms of the ubiquitous IQ test. The first such standardized testing was performed in 1904 by the psychologist Alfred Binet, who was commissioned by the French government to devise a method for identifying children whose sub-normal intellectual abilities (i.e., learning disabilities) might benefit from special schooling. Binet and his assistant Theophile Simon published their resulting test in 1905, which came to be known as the *Simon-Binet Scale* for assessment of mental ability (Strydom & Du Plessis, 2000). Lewis Terman, a psychology professor at Stanford University, later normalized the results by dividing the test subject's "mental age" (as reflected by the *Simon-Binet Scale*) by his or her chronological age, resulting in a so-called "intelligence quotient," or IQ. Terman published his results in 1916 as the *Stanford Revision of the Binet-Simon Scale of Intelligence*, generally abbreviated as *Stanford-Binet*, which became the standard IQ test in the United States for the next several decades (Linden & Linden, 1968).

By definition, IQ tests are intentionally structured to yield an average score of 100, with 50-percent of the population falling into the range of 90 to 110, leaving 25-percent above and the remaining 25-percent below. IQ scores tend to consistently rise a few points every 3 or 4 years, a phenomenon known as the "Flynn effect" (Flynn, 1994), and so tests must be periodically rewritten with harder questions to maintain a mean score of 100. Such standardized tests can be somewhat misleading indicators, however, providing only a general measure of academic intelligence. In fact, many argue they more aptly indicate what you have learned and how well you perform during testing, as opposed to your actual mental capacity. As William Calvin, a theoretical neurophysiologist at the University of Washington in Seattle, points out (Calvin, 1996):

"Intelligence gets framed in surprisingly narrow terms most of the time, as if it were some more-is-better number that could be assigned to a person in the manner of a batting average. It has always been measured by a varied series of glimpses of spatial abilities, verbal comprehension, word fluency, number facility, inductive reasoning, perceptual speed, deductive reasoning, rote memory, and the like."

Standardized tests also fall short in terms of assessing other important aspects of intelligence such as creativity and common sense. With no suitable definition for what intelligence actually is, some question the validity of an arbitrary scheme for its measurement. Indeed, I have known more than a few incredibly bright PhD types who clearly had way-above-average intelligence, yet possessed absolutely no practical knowledge or common sense. And what about the brilliant eccentrics whose private lives are in shambles because they lack the emotional intelligence to succeed in a relationship?

The Evolution of ROBART

Perhaps for lack of a better alternative, however, IQ scores continue to provide a rapid, fairly consistent, low-cost, and therefore widely accepted means of quantifying intelligence, sometimes with legal implications. Consider the case of convicted murderer Daryl R. Atkins, who along with an accomplice, kidnapped, robbed, and then killed a Virginia man in 1996. Atkins scored 59 on an IQ test two years later, well below the Virginia retardation threshold of 70. When reviewing his case in 2002, the U.S. Supreme court ruled execution of the mentally retarded as being unconstitutional. In a rather ironic twist of fate, however, the defendant's active participation in the intervening years of litigation seems to have had a positive effect on his mental capacities. His most recent test score was 76, now above the cut off. Consequently, Atkins is again on death row, and among the first of several such inmates awaiting trial to decide their mental status as a result of this ruling. Defense attorneys argue that: 1) IQ scores are rarely stable, 2) scores tend to rise over time, and, 3) practice can cause even further improvement (Union-Tribune, 2005).

But what does all this mean from a robotics perspective? Peter Van Turennout (1994), in his thesis dissertation entitled *Autonomous Motion on Wheels*, proposes the following definition of intelligence:

"The power of seeing, understanding, learning, knowing. Hence it implies the possession of knowledge, being able to manipulate and reason with this knowledge, the ability to gather new knowledge by observation and deduction, and to master the integration of all this."

I don't normally attach much significance to definitions per se, particularly with regard to complex and abstract concepts such as intelligence. But like the annoyingly invasive grain of sand that causes an oyster to ultimately produce a pearl, sometimes the end justifies the means, and we have to start somewhere. Much more could be said, obviously, but hopefully Van Turennout's grain-of-sand definition will be catalytic enough for our purposes here. In terms of artificial replication, we've done fairly well to date with his *seeing* and to some extent even *learning* criteria, but *understanding* and *knowing* remain difficult challenges.

19.1.2 Artificial Intelligence

The first book I ever read on this subject was entitled *The Handbook of Artificial Intelligence, Volume I*, a compilation of works from the preceding 25 years as edited by Avron Barr and Edward Feigenbaum (1981) of Stanford University. I was given this text by Dr. Jude Franklin, Director of the Navy Center for Applied Research in Artificial Intelligence, which had just been established at the Naval Research Laboratory in 1982. I visited the lab sometime in 1983, accompanied by RADM James Lisanby (SEA 90), on a quest to learn more about the Navy's interests in AI, and how it might potentially dovetail with NAVSEA's fledgling robotics program. My first and foremost question at the time (recall earlier discussion in Chapter 8) concerned the difference between AI and what I rather naively thought of as "normal" programming. The book provided the following description in its opening sentence:

Chapter 23 Intelligence Revisited

“Artificial Intelligence (AI) is the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior – understanding language, learning, reasoning, solving problems, and so on.”

Digging a little deeper, Jack Copeland (2000) lists three subcategories of AI: 1) *Strong AI*, 2) *Applied AI*, and 3) *Cognitive Simulation*. The first of these, *Strong AI*, is concerned with achieving human-equivalent performance on the part of an intelligent machine. The so-called *Strong AI Claim* asserts there is essentially no difference between a *living mind* and some inanimate *mind-like machine* or artifact that artificially emulates its functionality, a perspective which has come to be associated with a theory of mind known as *Functionalism* (Beloff, 2002). In other words, a computer-based instantiation of the human brain could theoretically be functionally equivalent to the original biological version. On the other hand, followers of *Applied AI* take a more pragmatic approach, pursuing less-ambitious bounded applications with some near-term commercial or military utility, and don’t get so wrapped up in the *Strong AI* debate. The third category of *Cognitive Simulation* is concerned with the use of computers to test candidate theories of how the mind performs certain tasks, such as problem solving or face recognition (Copeland, 2000).

The distinguished philosopher John Searle (2002) makes what I consider a very valid point in stating that the terminology “artificial intelligence” was probably less appropriate than “simulated cognition.” Certainly such a distinction would have cleared up some false expectations on my part in the early days with ROBART I. In fact, expectations for AI in general were way too high, even on the part of those more knowledgeable individuals directly involved in the research. And many of those not-so-directly-involved were quick to exploit the terminology anyway, often misusing it as an over-worked buzzword in countless robotic proposals that came across my desk in the mid-eighties timeframe, until it eventually became a joke in our office. Any tough engineering challenge not otherwise adequately addressed in the write up was simply to be solved “through the application of artificial intelligence techniques.” As a consequence, I distanced myself more and more from the terminology, and anyone who used it to excess without sufficient substantiation.

Enter Rodney Brooks of the MIT AI Lab, whom I had met through my association with Anita Flynn, as mentioned in Chapter 8. In what would come to be regarded as a watershed event, Brooks rocked the traditional AI community with the publication of “A Robust Layered Control System for a Mobile Robot” (Brooks, 1986), earning himself the prestigious title of “The Bad Boy of Robotics,” as bestowed by his ever-growing following of MIT grad students. The paper’s alternative slant pushed back against the prevailing conventional wisdom of the time (i.e., sense-think-act), in which *perception* was linked to *action* through *cognition*, as illustrated in Figure 19-2. Moravec’s Stanford Cart (Moravec, 1983; 1988) and SRI’s earlier robot *Shakey* (Nilson, 1984) effectively illustrate this classical approach, wherein a higher-level reasoning system operated upon the output of vision-based perception to provide the necessary motion planning for actuation (Brooks, 1999).

The Evolution of ROBART




Figure 19-2. The traditional AI approach linked *perception* to *action* through *cognition* (adapted from Brooks, 1999).

Brooks adopted instead a bottom-up behavior-based strategy (see Figure 19-3), proposing that perception and actuation were more directly linked, without the need for detailed intermediate world models relating one to the other. This simplistic reactive approach clearly resonated with me, since I had embraced the same philosophy on ROBART I (Everett, 1982).




Figure 19-3. Brooks proposed an alternative decomposition of the problem based upon task-achieving behaviors (adapted from Brooks, 1999).

Brooks went on to become one of the most respected and influential pioneers in the field, appointed Director of the MIT Artificial Intelligence Laboratory in 1997, and later promoted to Director of the combined Computer Science and Artificial Intelligence Laboratory in 2003. Outside of MIT (and a few other notable exceptions), however, progress for the most part was excruciatingly slow, with no significant successes. In the words of Jeff Hawkins, renowned inventor of the *Palm Pilot*, the field of AI rather lost its luster. Suffering from reduced funding, many researchers moved on to other fields, and more than a few start-up companies went out of business (Hawkins, 2002). In contrast, Brooks' Boston start-up, iRobot Corporation (which he co-founded with former MIT grad students Colin Angle and Helen Greiner), today mass-markets the behavior-based *Roomba* vacuum discussed earlier in Chapter 6, along with several other military and industrial systems.

Chapter 23 Intelligence Revisited

Towards the end of the century, after almost five decades of research, a rather sobering reality had settled in, as candidly expressed by Jack Copeland (2000):

“Excessive optimism in the 1950s and 1960s concerning Strong AI has given way to an appreciation of the extreme difficulty of the problem, which is possibly the hardest that science has ever undertaken.”

Computers still for the most part could not reliably perform many cognitive tasks normally associated with humans, such as understanding and/or translating language, interpreting visual imagery, simple abstract reasoning, and general problem solving.

Then into the picture steps (or rather jumps!) Ray Kurzweil (1999), a noted AI pioneer and successful entrepreneur with quite an impressive track record of technology predictions. While much of the rest of the learned population seemed convinced that computers would cease to function altogether come the new millennium, Kurzweil instead envisioned them becoming powerful enough to take over the world. All sensationalism aside, what Kurzweil did that impressed me was point out the more general application of *Moore's Law* across a broader domain than even Moore himself had envisioned, as we will examine below.

In 1965, Gordon E. Moore, then Director of Fairchild Semiconductor's Research and Development Laboratories, wrote a very insightful paper forecasting the monumental future of integrated circuits, with an emphasis on increased functionality and markedly reduced costs. In support of his more than optimistic predictions, Moore (1965) cited historical evidence from 1959 to 1965 regarding manufacturing costs per component as a function of chip complexity, which upon extrapolation to 1975 showed a very telling trend:

“The complexity for minimum-component costs has increased at a rate of roughly a factor of two per year (see graph on next page). Certainly over the short term this rate can be expected to continue, if not to increase. Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least ten years.”

Indeed, as President and CEO of Intel ten years later (1975), with more historical data available, Moore fine-tuned his prediction to the number of components doubling every 18 months, which became *Moore's Law* as we now generally know it (Schallur, 1996). In actuality, however, there is some disagreement within the semiconductor industry as to the exact timeframe for this doubling, varying from Moore's original prediction of one year, to as much as two and a half years (Geppert, 1998), with Moore himself now claiming two years (Kurzweil, 1999). Similarly, there is also considerable debate as to how long *Moore's Law* will continue to prevail, with some believing insurmountable barriers imposed by the physics of chip manufacturing will eventually flatten the curve and halt further progress (Schallur, 1996).

But as Kurzweil (1999) observes:

The Evolution of ROBART

“The speed and density of computing have been doubling every three years (at the beginning of the twentieth century) to one year (at the end of the twentieth century), regardless of the type of hardware used. Remarkably, this ‘Exponential Law of Computing’ has held true for at least a century, from the mechanical card-based electrical computing technology used in the 1890 US census, to the relay-based computers that cracked the Nazi Enigma code, to the vacuum-tube based computers of the 1950s, to the transistor-based machines of the 1960s, and to all of the generations of integrated circuits of the past four decades.”

Kurzweil points out that Hans Moravec (CMU) and David Waltz (NEC) also made the same observation during the 1980s timeframe. I could particularly relate, having employed vacuum-tube technology on *Walter*, electro-mechanical-relay logic for the *CRAWLER*, and finally moving into transistors and integrated circuits with the *ROBART* series.

Thus *Moore’s Law* was not the first paradigm, but instead the fifth in a continuing series to describe the exponential growth of computer technology, and it likely won’t be the last. This observation becomes Kurzweil’s foundation from which to extrapolate beyond the theoretical ceiling of transistors, and therefore say with confidence that computers will keep getting faster, achieving by 2020 a level of computational equivalence to the human brain. From that point onward, according to some, it’s all downhill for the human race as we know it today.

Perhaps at least part of the hype in Kurzweil’s presentation is intentional sensationalism for the sake of sales, and indeed the book has been very successful in that regard (i.e., on New York Times Bestseller List). What I find a bit disappointing is the fact that lots of revolutionary claims are made, but not always with adequate substantiation. Many of the critical (and currently unavailable) enabling technologies he cites depend heavily on the still evolving field of *nanotechnology*, in a manner suspiciously similar to the way “artificial intelligence techniques” were going to magically solve so many tough challenges in the past. *Nanotechnology*, for example, will supposedly provide the tiny robots (nano-bots) that will “non-invasively” travel through our blood vessels, reverse engineering our brains from the inside out, thereby allowing us to recreate them in silicon.

From my perspective, there are at least three areas here which warrant further consideration: 1) whether merely increasing computational capacity is sufficient guarantee that a computer can effectively surpass the capabilities of the human brain, 2) if futuristic *nano-bots* will be able to “non-invasively scan the brain” in any meaningful fashion to support its reverse engineering, and, 3) whether subsequent replication of a thusly scanned *human brain* would indeed capture the essence of the *human mind*.

For starters, increased computing power does not necessarily map over into a proportionally equivalent advancement in functionality, which Kurzweil readily admits. The mobility controller on ROBART III is perhaps a very good example. It began life

Chapter 23 Intelligence Revisited

(oops!) in 1992 as a 1-mHz 6502-based single-board computer with 32 kilobytes of RAM, hosting two Hewlett-Packard *HCTL-1000* PID controller chips on a daughter board. In late 1995 we upgraded to a 2-mHz 6502 to better enable dead-reckoning calculations. In 2000 one of my electrical engineers, Narek Pezeshkian, designed a 68HC11-based controller that had three onboard slots for *HCTL-1000* chips, but before we ever installed this version, we decided to upgrade further to a single-board *Power PC*, in the form of the *ipEngine* made by BrightStar Engineering. Then in 2004 we further evolved to a Compulab *686 CORE* with 128 megabytes of memory running at 266 MHz.

Yet throughout this continuing saga of evolutionary upgrades, the overall improvement in observed functionality (i.e., motor control, dead-reckoning calculations, collision avoidance) was rather minimal. What did substantially increase, unfortunately, was power consumption (by a factor of 12), and the likelihood of system failure. My engineering team, however, points out that the underlying motivation for this last upgrade was to support a standardized driver running under the USC open-source *Player/Stage* project, thus allowing all computer modules on the local area network to command actuators and retrieve sensor data using a TCP/IP interface.

With regard to the second point (i.e., the scanning nano-bots), as a practical roboticist I have some serious concerns with any such grandiose claims if not backed up by credible specifics. As I am often reminded these days by my friends in private industry, smaller is not necessarily easier or cheaper. We've had autonomous navigation and collision avoidance running on the golf-cart-sized MDARS-Exterior robots now since the mid-nineties (Heath-Pastore & Everett, 1994). But we are not yet able to effectively scale it down to the smaller man-portable robots like the iRobot *PackBot* or the Foster-Miller *Talon*, which are all still teleoperated (Everett, et al., 2004). So even further reduction to nano-scale systems that could maneuver in the blood stream, accurately perceive and interpret their surroundings, and globally keep track of their 3-D position and orientation with sufficient precision to map brain structure seems incredibly ambitious in comparison.

But let's just be charitably optimistic and assume for the moment such a nano-bot could in fact be created at some point in the future, perhaps by the very computers that will have by then exceeded my own mental limitations that now cause such skepticism. We are still faced with the venerable *asynchronous-data-registration problem*, as anyone who has ever written even primitive control software will readily attest. Suppose we were to request range data from a microprocessor that is pinging a multiple-transducer sonar array, as for example on *ROBART II* (Everett, 1985). The sonar driver runs in a continuous loop, timing sonar echoes on each pass, storing the resulting range values in memory for on-demand access by other processors in the system. But halfway through such a requested range-data retrieval (unless steps are taken to specifically preclude such), the sonar driver could conceivably overwrite the stored values with new ones from the most recent loop execution. The data actually transferred thus becomes a mix of old and new values representing two different sonar epochs, and thus is not an accurate portrayal of either.

The Evolution of ROBART

Now consider that the number of neurons in the human brain has been estimated at around a hundred billion, each of which is in contact with anywhere from one thousand to ten thousand of its neighbors. Consequently, the upper limit of potential connectivity permutations, or possible *brain states*, has been calculated to be in excess of the total number of atoms in the known universe (Ramachandran, 2004). So getting an accurate snapshot of brain connectivity is going to be rather problematic, to say the least. A single nano-bot cruising this Disneyland of synaptic wonderment is certainly not going to cut it. At the other extreme, if we assume a dedicated nano-bot for each possible permutation, the subject's head will explode long before the equivalent mass of our entire universe can be suitably emplaced within the blood vessels of his or her brain. Hardly what I'd call non-invasive.

So let's say we were to pick a happy medium somewhere in between these two extremes (maybe just a migraine-inducing number of scanner-bots), since we don't really have to cover all the possible permutations, just the established ones at some particular point in time. (And we'll also ignore the possibility that these millions of cruising nano-bots may log-jam in the narrowest reaches of our cranial capillaries, spawning a stroke-inducing artificial blood clot.) Assuming the ambitious scanners all worked as advertised, how would we synchronize the resulting data flow in order to accurately reconstruct it? For that matter, how would we actually *get* the data, given that effective communication pathways are already the major bottleneck of even today's more primitive computers?

It's been my observation over the years that when it comes to robotic projections, the degree of expressed optimism is pretty much inversely proportional to the amount of actual hands-on experience. In any event, all such speculation is not the relevant issue here, as conceivably man ultimately *will* find a way to at least partially reverse-engineer the brain, even if scanning nano-bots in the bloodstream never come to pass. The real issue to me is point number three: whether or not such an incredible feat of replication would actually yield an artifact fully representative of the human mind.

Consider the popular analogy that the *mind* is to the *brain* as computer *software* is to computer *hardware*. Now suppose aliens (as a convenient third party) wanted to reverse-engineer some strange computer they had salvaged from a NASA space probe that crash-landed on their planet. Lacking any supporting information on what this computer did or how it did it, alien engineers might conceivably recreate the *hardware* with an exact functional equivalent, using their own (presumably more advanced) technology. In addition, if the *software* that had been running on this computer was embedded as logical ones and zeroes within non-volatile hardware memory, then it could likewise be recovered.

That being the case, the entire system (hardware and software) could theoretically be reinstituted in an altogether different form, but with completely identical function. Furthermore, this functional restoration could be done even if the ambitious aliens never fully understood how the original system worked at any point in the entire process. Following analogous reasoning, the supporters of *Strong AI* believe a non-invasive brain

Chapter 23 Intelligence Revisited

scan (or some similar reverse-engineering strategy) will someday enable us to download our minds into thinking machines, even if we don't understand what the mind actually does or how it does it.

There's also the question of whether a single snapshot of the brain would adequately describe its behavior over time, taking into account such things as mood swings, for example. Whitby (2003) states: "One of the most important [differences] is that brains are immersed in a complex and changing mix of chemicals which constantly affects their performance." Through a chemical process known as neuromodulation, the brain is able to change its state on a global basis (Hobson, 2002). As an example, the transition from awake to asleep is chemically induced, wherein one part of our brain alters the chemistry of another part, causing us to lose consciousness through a process of cortical inhibition (Stark, 1993). When the brain self-activates later during REM sleep, its chemical state changes yet again.

Along the same lines, increased levels of the hormone testosterone have been shown to enhance the probability of aggressive behavior, while marijuana usage tends to have a somewhat opposite effect. A recent study by Swiss and American researchers suggests the neuromodulator *oxytocin*, a natural hormone released in the brain during eating, sex, childbirth, and lactation, is directly linked to trust (Kosfield, et al., 2005). Inhaling a synthesized form of *oxytocin* (normally used to stimulate lactation or induce contractions during labor) made study participants 17-percent more likely to trust someone else with their money.

Clearly the chemical-biasing effect of neuromodulation adds an entirely new and temporal dimension to classic synaptic transmission (Greenfield, 1997), and the resulting number of connectivity permutations (i.e., possible brain states) is therefore essentially limitless. Accordingly, we cannot just capture the brain's *static connectivity status* at some particular point in time, because the *effective connectivity* undergoes significant temporal variation as a result of this neuromodulation. What happens to the master schematic as seen by the probing nano-bots when some external event causes us to switch suddenly from upbeat and optimistic to depressed and cynical? What if we were to scan a schizophrenic? Or suppose we replicate the brain based on a one-time scan of an enthusiastic happy person; would not the results be different than if we scanned this same individual at another time when they were clinically depressed? Would the new brain then retain this gloomy state of depression as a permanent feature (i.e., bitch-bot from hell)? Or is the mechanism that regulates moods a part of that which gets scanned, and hence still performing its role in the artifact?

In addition to moods, consider also emotions. Would a robot with an artificial brain be capable of the full range of human emotions, triggered by time-variant stimuli? For example, what exactly makes us recognize and react to humorous situations? Can an artifact likewise have a sense of humor, maybe even crack a joke? For what purpose do humans even have a sense of humor? And what, if any, is the evolutionary function of sarcasm?

19.1.3 Artificial Common Sense

But does a robot necessarily need to fully emulate our full spectrum of natural intelligence, mood swings and all? Certainly there are plenty of bona fide applications, some already in practice, where such sophistication is not necessary. (Indeed, one of the main advantages of a robot in certain applications is the very fact that they are *not* emotional, and therefore not averse to dangerous missions.) My goal as a roboticist all these years has not been to create an artificial human per se, but rather to produce a robotic system capable of sufficiently sophisticated behavior to where it can intelligently interact with humans as part of a synergistic team. If the resulting machine behaviors are very human-like, maybe even human-equivalent in some cases, then so be it. But I don't fixate on a holy grail of cloning human intelligence just for the sake of doing it.

Rudimentary common sense, on the other hand, might indeed be a worthy pursuit. Some overarching awareness of what's actually going on certainly helps humans make better decisions when confronted with an unexpected occurrence, and lack of such on the part of a robot could be a critical deficiency in some applications. Would you, for example, trust a baby-sitter that had no real concept of the dangers associated with caring for small children? Envision a futuristic "robo-nanny," carefully pre-programmed with a number of suitable if-then rules related to child care: *1) if baby hungry, then give food, 2) if diaper dirty, then change diaper; 3) if baby tired, then put to bed*, and so forth. But what if the house were to catch fire, and no appropriate if-then rule specifically applied? The robot, with no *conscious awareness* of the danger, fails to act accordingly. It might try to extinguish the fire, without first moving the baby to safety.

But just for the sake of argument, let's say some astute programmer had the foresight to anticipate such a problem, and added: *4) if smoke detector alarms, evacuate house*. That being the case, let's now consider a slightly revised scenario, where the robot looks out the window on an otherwise uneventful day and perceives that the next-door neighbor's house is on fire. Beyond it several more houses and the surrounding landscape are burning, the flames fanned high by strong Santa Ana winds blowing in from the California desert. The family dog is whining, nervously pacing back and forth, causing the baby to cry. The all important smoke detector, however, remains silent. The nanny-bot assumes the crying baby is tired and puts it to bed, oblivious to the impending danger. By the time the smoke detector does respond, it may be entirely too late for escape.

The point I'm getting at here is that humans have an inherent *awareness* of what's happening around them, allowing them to evaluate and assess things in context and make predictions about what may happen next. This learned capacity for prediction extends the intelligent behavior of a human to a level above that of an otherwise intelligent robot lacking any such comparable awareness. Indeed, Hawkins (2004) specifically proposes that this very ability to predict is the more appropriate *Turing Test*, as opposed to just the appearance of intelligent behavior.

"Prediction is not just one of the things your brain does. It is the *primary function* of the neocortex, and the foundation of intelligence. The cortex is an organ of prediction. If we want to understand what intelligence is, what creativity is, how the brain works, and how to

Chapter 23 Intelligence Revisited

build intelligent machines, we must understand the nature of these predictions and how the cortex makes them. Even behavior is best understood as a by-product of prediction.”

But while it may indeed be possible to create some partial capability of inference along these lines, to date it has only been done within very narrow domains of specialty (for example, expert systems). Furthermore, such limited powers of inference are by no means comparable to true awareness. No robot has today, nor probably will have in the foreseeable future, what can unequivocally be labeled as *common sense*. And Copeland (2000), whom I’d classify as a proponent of *Strong AI*, specifically states that even “expert systems have no common sense.”

But what exactly does it mean for a computer program to have common sense? John Beloff (2002) lists the following criteria for what “something” must do, above and beyond just the mechanical processing of information (i.e., like a computer), in order to say it was actually *thinking*: 1) be *aware* of what it is doing, 2) know what it is thinking about, and, 3) recognize when it has reached a conclusion. In other words, he summarizes: “it should have insight into what is going on.”

19.2 The Meaning of Consciousness

Such an “insight into what is going on,” as Beloff put it, implies *conscious awareness*. The issue of consciousness has been debated for centuries, and a plethora of contemporary writings describing various theories abound (i.e., Blackmore (2004); Chalmers, (1996); Damasio (1999); Dennett (1991); Edelman & Tononi (2000); LeDoux, (2002); Pinker (1997); Ramachandran (2004); Searle (1997); and many others). But what exactly does it mean to be “conscious” of something, and how would we define such a thing, in order to better assess the likelihood of its artificial replication? The simple answer is we can’t, yet. Churchland (1995) sums it up nicely: “Definitions are best framed after we have settled on an adequate understanding of what needs defining. And that is something we won’t have until we possess an adequate scientific theory of consciousness.”

What is doable, however, is an articulation of that which we do know (or have at least perceived) in the way of the salient features of consciousness, for which Churchland offers up the following:

Involves short-term memory – Consciousness typically exhibits some grasp of a temporal relationship with regard to our actions, in terms of how our present relates to our past, and projections for its future.

Independent of sensory inputs – Consciousness is not dependent upon external stimuli, in that we can be consciously aware of our thoughts, memories, and other subjective internal states, even while blocking all outside sensory experiences.

Displays steerable attention – We can selectively focus our conscious attention at will, thus filtering out distractions, allowing the brain to concentrate on the higher priority concerns.

The Evolution of ROBART

Allows alternative interpretations – Consciousness provides for the generation and subsequent assessment of competing interpretations in response to complex, ambiguous, or incomplete sensory input.

Disappears in deep sleep – We lose consciousness altogether (see next section) upon entering a state of deep sleep.

Reappears in dreaming – Some modified or muted form of consciousness seems to re-emerge during our dreams, while we are still asleep, allowing us to experience and even remember a dream upon awakening.

Provides a single unified experience – The cognitive inputs from our various sensor modalities are blended into an integrated representation of the collective experience, rather than individually generating separate and uncorrelated awarenesses of different aspects of the same experience. We perceive a *red bicycle*, for example, even though the representation for *shape* is processed and maintained in a different part of our brain than the representation of *color*.

Churchland goes on to suggest how “a suitably configured recurrent network will display cognitive behaviors that are systemic functional analogs of all seven of these familiar dimensions of consciousness.” In other words, if all these salient features of consciousness can be artificially demonstrated on an individual basis, then what is theoretically stopping the sum of the parts?

V.S. Ramachandran, Director of the Center for Brain and Cognition, and Professor of Psychology and Neuroscience at the University of California, San Diego, provides a similar listing from a slightly different perspective, based upon his belief that *consciousness* boils down to two mutually dependent issues: 1) the age-old problem of *self* (i.e., our individualistic existence as a unique entity), and 2) the more recently introduced problem of *qualia* – subjective sensations: the what-it's-like character of, or way it feels to have, mental states such as pain, seeing red, smelling a rose, etc. (Eliasmith, 2004). Their interdependence lies in Ramachandran's observation that experienced sensations (i.e., *qualia*) inherently require someone to experience them, just as a *self* cannot be devoid of sensory experiences and emotions. He lists five defining characteristics of *self* (Ramachandran, 2004):

Continuity – An unbroken temporal thread that links all aspects of existence (i.e., experience) in terms of a past, a present, and a future.

Unity – Experiencing this existence from the single perspective of an individual, unique among many (all with their widely diverse experiences, thoughts, memories, beliefs, and subsequent cultures).

Embodiment – A sense of being anchored to our physical bodies.

Agency – A sense of free will, wherein we are in charge of our own actions and destinies.

Reflection – An introspective awareness of ourselves.

In summary, the fundamental question of the *Strong AI Claim* reduces to one of the *artificial self*, and whether such (as described above) can be created in a machine, thus making it *conscious* of its existence, *reflectively aware* of its *current* and *past* actions,

and ultimately attaining the *free will* to govern its own *future*. In other words, can a machine have a feeling of identity (i.e., an *artificial self*) and be consciously aware of its perceptions and/or actions? So far the issue remains far from resolved.

19.3 *Unconscious Versus Subconscious*

Before much further contemplation of this *artificial conscious awareness*, we should probably clear up one source of potential confusion that arises from how the terminology is applied in describing our own mental states as human beings. The term *unconscious* obviously implies the opposite state of *conscious*. But when used in the one sense, *unconscious* refers to a state of markedly reduced brain activity, such as when a person is “knocked unconscious” by a blow to the head, is under anesthesia, or is in a coma. On the other hand, many authors I have cited use the term *unconscious* in the same way I have been using the term *subconscious*: to denote a state of considerable brain activity, but one which differs from the *conscious state* in that we are not acutely aware of it.

To help minimize further confusion, I shall continue to use the terminology *subconscious mind* as opposed to the synonymous *unconscious mind*, as to me the term *unconscious* implies an abnormal state (i.e., knocked out, in a coma, intoxicated, or drugged). Some interesting questions arise when contemplating this confusing terminology, however. When we are *unconscious*, are we truly not aware? Can a patient in a coma possibly hear things being said by the surrounding medical staff? There is evidence that some patients can recall (under hypnosis) an almost verbatim account of surgical-team dialogue that took place while they were under general anesthesia (Ratey, 2001). There are also a number of intermediate states between *conscious* and *unconscious*, such as *sleep*, *somnolence*, *stupor*, and *vegetative*, but to keep from diverging too far from the real point, we shall consider only the first of these.

It seems pretty clear that being *asleep* is quite different from being *unconscious* as defined earlier. For one thing, our brains are almost as active while we are sleeping as when we are awake. Certainly we all have thought processes that occur during sleep, many of which we remember later when we wake and recall our dreams. Some of these thought processes actually cause physical reactions even as we sleep, such as tossing and turning, grinding of teeth, talking, perhaps even the extreme of getting up out of bed and sleepwalking. On the other hand, our muscles are inhibited during REM sleep to disable the motor system, and any sensation of movement during dreaming is purely imagined. This chemically induced paralysis may help explain that helpless sensation of running in slow motion to escape an imaginary villain, which so often occurs in nightmares.

What about Paul Churchland’s theory of some “muted state of consciousness” during dream sleep? Some people, technically termed *oneironauts*, indeed seem very self-aware under conditions known as *lucid dreaming* (LaBerge, 1991). Such individuals are not only conscious of the fact they are dreaming, but even more astounding, are actually able to control how the dream unfolds. This is perhaps the ultimate *virtual reality* experience for those that can do it, as they are reportedly able to dream about anything they wish, how they wish, and will even remember it later when they wake up! Why go rent a movie? Unfortunately, it supposedly is not easy to learn how to do such a thing.

The Evolution of ROBART

But seriously, if we accept the broad definition of dreaming as mental activity occurring in sleep, it seems logical that some degree of consciousness be associated with that activity. Indeed, J. Allan Hobson, Professor of Psychiatry at Harvard Medical School and one of the leading researchers in the field, states that “dreaming is a distinctive form of *conscious awareness* caused by the state of the brain in sleep” (Hobson, 2002). Hobson lists the following cognitive features of dreaming: 1) loss of self awareness, 2) loss of orientational stability, 3) loss of directed thought, 4) reduction in logical reasoning, and 5) poor memory during and after the dream.

From Hobson’s perspective, waking and dreaming are two distinct states of consciousness, with their differences determined by brain chemistry:

“To cut a long story short, the serotonin and noradrenaline cells that modulate the brain during waking reduce their output by half during non-REM sleep, but are shut off completely during REM sleep. This means that the electrically reactivated brain is working without the participation of two of its chemical systems that mediate the awake state. These very systems have been strongly implicated in precisely those awake-state functions (such as attention, memory, and reflective thought) that are lost in dreaming.... The main point to keep in mind is that in REM sleep the brain, although electrically activated as in waking, is activated in a chemically very different way.”

It’s interesting to note that upon awakening from *normal* sleep, we usually have a fairly accurate sense of time passage, often knowing within a few minutes or so what time it is. (I stress normal, because if our sleep pattern changes from the norm, as for example due to jet lag, we can sometimes awaken quite confused as to even what day it is.) On the other hand, anyone who has ever been anesthetized for surgery (i.e., rendered unconscious) typically recalls awakening in what seems like a matter of seconds, even though the procedure may have taken several hours or more. This undistorted sense of time passage during sleep further underscores the fact that our *subconscious mind* is working away as we doze, taking care of business, and keeping track of the time as it does so.

But what business would that be? It is generally viewed that sleep is a regenerative process for the body in general, and the brain is no exception. Our everyday experiences make it pretty clear that sleep deprivation leads to mental exhaustion, impairing our ability to think and reason clearly. But there are still a lot of unanswered questions as to why this happens. Some believe that sleep provides an opportunity for the brain to replenish the chemical neurotransmitters that enable the passage of electrical signals from axons to dendrites in our synapses. The signal is not transmitted through direct electrical contact, but by the electrically stimulated chemical release of these neurotransmitters into the synaptic cleft. The infusion of this small amount of fluid from the axon changes the chemical properties within the synapse, thus influencing the electrical charge of receptors on the receiving dendritic surface (Searle, 1997). If the source chemicals within the neuron are depleted, it follows that our normal brain activity would suffer accordingly.

Chapter 23 Intelligence Revisited

Many researchers have suggested that one of the biggest functions performed during sleep is the consolidation and optimization of memories (Hobson, 2002). We take in an enormous amount of asynchronous information during the course of our day, much of which is extraneous and even distracting from the really important issues. During sleep this data-acquisition is temporarily suspended, giving the *subconscious mind* time to sort things out and clean up our “scratch-pad” memory for further use. We wake up the following morning with a clear head, ready to take on new challenges. (For me, two other activities that seem to have a similar “head-clearing” effect are taking either a long walk or a hot shower.) Relevant information considered potentially useful is transferred to *long-term memory*, where it is filed away in a more organized fashion that facilitates later retrieval. Equally important, this reorganization is also believed to assist in the all-important association of related issues, a key part of the so-called “binding problem.”

Consider the following analogy. Despite the well-intentioned efforts of our high-school literature teachers, it’s been my observation that most engineers typically do not begin drafting a journal or conference paper with an organized outline of the relevant issues. Instead, we lean more towards a bottom-up approach: jotting down anything even remotely relevant to the subject, usually with little regard to coherent flow, sometimes not even sure yet of the main point we ultimately wish to convey. Later, through an iterative process of incremental improvement, we massage this rough draft into some semblance of order, grouping together related thoughts, tossing out the less significant details to meet the page-count restriction, and hopefully concluding with some reasonable meaning for it all.

As we sleep, our *subconscious mind* goes through similar motions in sorting out the day’s incredible hodge-podge of wide-ranging experiential input, building associations among related concepts, then drawing inferences from this more organized synergistic representation. As a result, we often wake up with a much clearer perspective, sometimes even creative insights. Another way of looking at it would be the analogy of sifting through one’s “in-box” at the end of a hectic workday, tossing out the junk-mail, then alphabetically filing any important documents by subject in appropriate folders for future reference. We certainly seem aware of the effect, if not the process specifics, in that we often intentionally defer important decisions until such time as our head clears, perhaps even saying, “Let me sleep on it.”

There is also strong evidence this “sorting-out” procedure is not limited to just our mental representations, but applies as well to the learning of coordinated motor skills. For example, I used to water-ski along the coastal tidewaters of South Carolina in my youth, on almost a daily basis during the summer months. When I left the region later in life, my opportunities to ski became more and more infrequent, generally limited to about a week or so each summer when I went home to visit. On such occasions I would invariably start off quite rusty, to where it usually took the full week to regain my former prowess, but by that point it was sadly time to leave.

As the years went by I began to notice that my improvement in performance was always most apparent the following day, whereas diminishing returns were achieved by

putting in extended hours all at one time. I remember thinking some “offline post-processing activity” seemed to be optimizing my motor-reflex coordination overnight, but I had no idea as to what or how. Hobson (2002) indicates this optimization process may take even longer, several days or perhaps even a week.

19.4 Summary

I’m sure much of my admitted ignorance with respect to brain function can be readily explained by specialists in the appropriate fields, and answering all the hypothetical questions posed in this chapter is neither within my current capabilities nor even my purpose in bringing them up. My intent is simply to point out that the brain is not a static configuration of neural gates that can be scanned at some specific instant in time for subsequent replication, but rather a very dynamic system capable of localized as well as global state changes that are chemically induced. Furthermore, these changes do not occur as instantaneous step functions, as Hobson (2002) points out:

“As the state of the brain changes continuously, it only gradually changes its mode. It does not suddenly switch from one state to another. Furthermore, the continuous and gradual modulatory changes do not affect every single neuron of the brain identically or even simultaneously. These generalizations, which all flow from neurobiological work on the state control systems of the brain stem, have far reaching implications for a general theory of mind as well as for a specific theory of dreaming.”

I think it is also important to understand that the brain is not a stand-alone entity to be copied as such, but operates in conjunction with the entire body, the two being inextricably intertwined as a fully integrated system. This is not a symbiotic relationship but an absolutely essential one, for neither component has even minimal utility without the other (i.e., a brain without a body is no less dysfunctional than a body without a brain). For that reason, merely scanning a brain for artificial replication seems like a gross oversimplification of reality.

If we truly aspire to emulate the phenomenal capabilities of the human brain, we need to look closely at the bigger picture, as there’s a lot more happening in our minds than we can currently explain. Even if we ignore the *sleep state* and the *unconscious state*, and consider only the *waking state*, we have at least two coexisting but decidedly different thought processes to take into account here, both the *conscious* and the *subconscious*. As we shall explore further in the next chapter, what we don’t yet realize about their complex interaction may have noteworthy ramifications further down the proverbial road. Hawkins (2004) sums it up thusly:

“AI scientists tried to program computers to act like humans without first answering what intelligence is and what it means to understand. They left out the most important part of building intelligent machines, the intelligence! ‘Real intelligence’ makes the point that before we attempt to build intelligent machines, we have to first understand how the brain thinks, and there is nothing artificial about that. Only then can we ask how we can build intelligent machines.”

19.5 References

- Armstrong, D.M., "What is Consciousness," *Nature of Mind and Other Essays*, Cornell University Press, Ithaca, NY, 1981.
- Barr, Avron, and Feigenbaum, Edward A., eds., *The Handbook of Artificial Intelligence, Volume I*, ISBN 0-86576-005-5, HeurisTech Press, William Kaufman, Los Altos, CA, 1981.
- Beloff, John, "Minds or Machines," *Truth Journal*, Leadership U, 2002.
- Blackmore, Susan, *Consciousness: An Introduction*, ISBN 0-19-515342-1, Oxford University Press, Oxford, 2004
- Brooks, Rodney A., "A Robust Layered Control System for a Mobile Robot," *IEEE Journal of Robotics and Automation*, RA-2, April, 1986.
- Brooks, Rodney A., *Cambrian Intelligence: The Early History of the New AI*, ISBN 0-262-02468-3, MIT Press, Cambridge, MA, 1999.
- Brooks, Rodney A., *Flesh and Machines: How Robots Will Change Us*, ISBN 0-375-42079-7, Pantheon Books, Random House, New York, NY, 2002.
- Calvin, William H., *How Brains Think: Evolving Intelligence, Then and Now*, ISBN 0-465-07278-X, BasicBooks (Perseus Books Group), New York, NY, 1996.
- Churchland, Paul M., *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain*, ISBN 0-262-03224-4, MIT Press, Cambridge, MA, 1995.
- Copeland, Jack, *Artificial Intelligence: A Philosophical Introduction*, ISBN 0-631-18385-X, Basil Blackwell, Oxford, 1993.
- Copeland, Jack, "What is Artificial Intelligence?" (http://www.cs.usfca.edu/www.AlanTuring.net/turing_archive/index.html), May, 2000.
- Damasio, Antonio, *The Feeling of What Happens*, ISBN 0-15-601075-5, Harvest (Harcourt), San Diego, 1999.
- Dawkins, Richard, *The Selfish Gene*, ISBN 0192860925, Oxford University Press, Oxford, 1990.
- Dennett, Daniel C., *Consciousness Explained*, ISBN 0-316-18065-1, Little, Brown and Co., Boston, 1991,
- Edelman, Gerald, and Tononi, Giulio, *A Universe of Consciousness: How Matter Becomes Imagination*, ISBN 0-465-01377-5, Basic Books, New York, 2002.
- Eliasmith, Chris, ed., *Dictionary of Philosophy of Mind*, on the specific topic of "qualia," <http://www.artsci.wustl.edu/~philos/MindDict/index.html>, 2004.
- Everett, H.R., *A Microprocessor Controlled Autonomous Sentry Robot*, Masters Thesis, Naval Postgraduate School, Monterey, CA, October, 1982.
- Everett, H.R., "A Multi-Element Ultrasonic Ranging Array," *Robotics Age*, pp. 13-20, July, 1985.
- Everett, H.R., and Pacis, Estrellina, "Towards a Warfighter's Associate: Eliminating the Operator Control Unit," SPIE Mobile Robots XVII, International Symposium on Optics East, Philadelphia, PA, 25-28 October, 2004.
- Flynn, James R., "IQ Gains Over Time," in *Encyclopedia of Human Intelligence*, R.J. Sternberg, ed., MacMillan, New York, pp. 617-623, 1994.
- Geppert, Linda, "The Media Event: Moore's Law Mania," *IEEE Spectrum*, Vol. 35, No. 1, January, 1998.

The Evolution of ROBART

- Greenfield, Susan A., *The Human Brain: A Guided Tour*, ISBN 0-465-00726-0, Basic Books, New York, NY, 1997.
- Hawkins, Jeff, and Blakeslee, Sandra, *On Intelligence*, ISBN 0-8050-7456-2, Times Books, Henry Holt and Company, New York, NY, 2004.
- Heath-Pastore, T.A., Everett, H.R., "Coordinated Control of Interior and Exterior Autonomous Platforms," ISRAM 94, Fifth International Symposium on Robotics and Manufacturing, Maui, HI, August, 1994.
- Hobson, J. Allan, *Dreaming: An Introduction to the Science of Sleep*, ISBN 0-19-280304-2, Oxford University Press, Oxford, 2002.
- Hunt, Morton, *The Universe Within: A New Science Explores the Human Mind*, ISBN 0-671-25258-5, Simon and Schuster, New York, NY, 1982.
- Johnson, Steven, *Mind Wide Open: Your Brain and the Neuroscience of Everyday Life*, ISBN 0-7432-4165-7, Scribner, New York, NY, 2004.
- Kosfeld, M., Heindrichs, M., Zak, P.J., Fischbacher, U., and Fehr, E, *Nature*, **435**, Pp. 673-676, June 2, 2005.
- Kurzweil, Ray, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, ISBN 0-670-88217-8, Penguin Books, New York, NY, 1999.
- Kurzweil, Ray, "The Evolution of Mind in the Twenty-First Century," in *Are We Spiritual Machines? Ray Kurzweil vs. the Critics of Strong AI*, Richards, Jay W., ed., ISBN 0-963854-3-9, Discovery Institute, Seattle, WA, 2002.
- LaBerge, Stephen, *Exploring the World of Lucid Dreaming*, ISBN 0-34-537410-X, Ballantine Books, 1991.
- LeDoux, Joseph, *Synaptic Self: How our Brains Become Who We Are*, ISBN 0-670-03028-7, Penguin Books, New York, 2002.
- Linden, K.W., and Linden, J.D., *Modern Mental Measurement : A Historical Perspective*, Houghton Mifflin, Boston, MA, 1968.
- Lycan, William G., and Zena C. Ryder, "The Loneliness of the Long-Distance Truck Driver," *Analysis* 63, pp. 132-36, 2003.
- Moore, Gordon E., "Cramming More Components Onto Integrated Circuits," *Electronics*, Vol. 38, No. 8, April 19, 1965.
- Moravec, Hans P., "The Stanford Cart and the CMU Rover," IEEE [?], 1983.
- Moravec, Hans, *Mind Children: The Future of Robot and Human Intelligence*, Harvard University Press, Cambridge, MA, 1988.
- Moravec, Hans, *Robot: Mere Machine to Transcendent Mind*, ISBN 0-19-511630-6, Oxford University Press, Oxford, 1999.
- Nilson, J.J., *Shakey the Robot*, Technical Note 323, SRI AI Center, Menlo Park, Calif., 1984.
- Pacis, Estrellina, Everett, H.R., and Dave Bruemmer, "Enhancing Functionality and Autonomy in Man-Portable Robots," Proceedings, SPIE Unmanned Ground Vehicle Technology VI, Defense and Security, Orlando, FL, 12-16 April, 2004.
- Pinker, Steven, *How the Mind Works*, ISBN 0-393-04535-8, W.W. Norton, New York, 1997.
- Ramachandran, V.S., *A Brief Tour of Human Consciousness*, ISBN 0-13-148686-1, Pi Press, Pearson Education, New York, 2004.

Chapter 23 Intelligence Revisited

- Ratey, John J., *A User's Guide to the Brain: Perception, Attention, and the Four Theaters of the Brain*, ISBN 0-679-45309-1, Pantheon Books, Random House, New York, 2001.
- Schallur, B., *The Origin, Nature, and Implications of Moore's Law: The Benchmark of Progress in Semiconductor Electronics*, available online at <http://mason.gmu.edu/~rschalle/moorelaw.html>, September 26, 1996.
- Searle, John R., *The Mystery of Consciousness*, ISBN 0-940322-06-4, New York Review, New York, 1997.
- Searle, John, "I Married a Computer," *Are We Spiritual Machines? Ray Kurzweil vs. the Critics of Strong AI*, Richards, Jay W., ed., ISBN 0-963854-3-9, Discovery Institute, Seattle, WA, p. 69, 2002.
- Sharples, Mike, Hogg, David, Hutchison, Chris, Torrance, Steve, and David Young, *Computers and Thought: A Practical Introduction to Artificial Intelligence*, Cogswell Project, October 22, 1996.
- Simons, Geoff, *Robots: The Quest for Living Machines*, ISBN 0-304-34086-3, Cassel, London, 1992.
- Stark, Todd I., "Hypnosis.com FAQ," www.hypnosis.com/faq/, July, 1993.
- Strydom, J., Du Plessis, S., "IQ Test: Where Does It Come From and What Does It Measure?," http://www.audioblox2000.com/dyslexia_dyslexic/dyslexia014.htm.
- Union-Tribune, "Death Row Inmate's Fate Could Hinge on his Rising IQ Score," San Diego Union-Tribune, p. A21, February 6, 2005.
- Van Turennot, Peter, *Autonomous Motion on Wheels*, ISBN 90-9006530-X, Ph.D. Thesis, Universiteitsdrukkerij T.U. Delft, 1994.
- Warwick, Kevin, *March of the Machines: The Breakthrough in Artificial Intelligence*, ISBN 0-252-07223-5, University of Illinois Press, Chicago, 1997.
- Winston, Patrick Henry, *Artificial Intelligence*, ISBN: 0201084546, Addison-Wesley, 1977.
- Whitby, Blay, *Artificial Intelligence: A Beginner's Guide*, ISBN 1-85168-322-4, Newworld Publications, Oxford, 2003.